

Surrogate-Assisted Multiobjective Co-Design of Fractional-Order Aerospace DC–DC Converters under Thermal, Transient, and Fault-Survivability Constraints

Taylan Demir

1996.matematik.1996@gmail.com

Ankara University <https://orcid.org/0000-0001-7402-4468>

Atakan koçyiğit

Atılım University <https://orcid.org/0009-0005-6060-0547>

Research Article

Keywords: fractional-order control, DC–DC converter, more-electric aircraft, multiobjective optimization, surrogate modeling, machine learning, thermal constraint, fault ride-through, power electronics co-design

Posted Date: September 17th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11066555/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Surrogate-Assisted Multiobjective Co-Design of Fractional-Order Aerospace DC–DC Converters under Thermal, Transient, and Fault-Survivability Constraints

Taylan Demir^{1,*}

Atakan Koçyiğit²

¹Department of Mathematics, Faculty of Science, Ankara University, Ankara, Türkiye

²Department of Electric and Electronic Engineering, Faculty of Engineering, Atılım University, Ankara, Türkiye

*Corresponding author: demir.taylan96@gmail.com

Abstract

More-electric aircraft power converters are designed under competing requirements rather than a single performance target: rapid bus-voltage recovery, low semiconductor loss, moderate passive-component size, thermal headroom, and fault ride-through must be achieved simultaneously. This paper formulates these requirements as a six-variable co-design problem for a three-phase interleaved 270 V-to-28 V, 800 W DC–DC converter controlled by a fractional-order PI^λ law. The decision vector contains switching frequency, phase inductance, output capacitance, fractional integral order, and two controller gains. A physics-based averaged model supplies the training data, while a machine-learning layer is used only to accelerate design-space screening. Five thousand Latin-hypercube designs are evaluated with the nonlinear averaged model; 3595 satisfy the imposed dynamic and hardware constraints. An ensemble feasibility classifier attains 96.4% test accuracy, 99.0% precision, and 96.0% recall. A multilayer surrogate for the transient objective has a mean absolute percentage error of 7.18% and a median error of 3.56%; because its accuracy deteriorates near saturation boundaries, every reported Pareto design is re-evaluated with the original physics model. Screening 60,000 candidates yields 26,209 high-confidence feasible points, from which 68 trade-off designs are retained by a bank of Dirichlet scalarizations and subsequently verified. A balanced design at $f_s = 246.23$ kHz, $L = 22.53$ μ H, $C = 960.91$ μ F, and $\lambda = 0.9042$ reduces the composite transient index by 13.8%, converter loss by 10.0%, and simulated fault sag by 36.2% relative to a conventional baseline. Its conservative capacitor ride-through allowance rises from 0.150 ms to 0.240 ms, while the normalized passive sizing index decreases slightly by 2.1%. The price is a larger phase-current ripple, which remains below the imposed 8 A limit. The study therefore treats machine learning as a screening instrument, not as a replacement for converter physics, and exposes the design compromises that are hidden by single-objective tuning.

Keywords: fractional-order control; DC–DC converter; more-electric aircraft; multiobjective optimization; surrogate modeling; machine learning; thermal constraint; fault ride-through; power electronics co-design.

1 Introduction

The electrical architecture of more-electric and all-electric aircraft is moving toward higher DC bus voltages, distributed conversion stages, and increasingly software-defined power management. These changes improve wiring mass and controllability, but they also transfer a larger fraction of system-level reliability to power-electronic converters and their supervisory control. Recent work on future aircraft DC systems emphasizes reliability, fault management, electromagnetic compatibility, power quality, and converter efficiency as simultaneous design requirements rather than separable afterthoughts [11, 12, 14, 13]. The difficulty is structural: increasing switching frequency may reduce passive-component size yet increase switching loss and thermal stress; enlarging the output

capacitor increases fault hold-up time but adds volume; aggressive feedback gains improve transient recovery but may amplify ripple, excite unmodeled dynamics, or consume duty-cycle margin.

Wide-bandgap semiconductors have enlarged the practical design space. SiC and GaN devices permit higher switching frequencies and higher power density than conventional silicon devices, while creating new thermal, packaging, reliability, and EMI constraints [8, 9, 10]. Consequently, converter design is no longer well represented by a sequence in which the power stage is fixed first and the controller is tuned afterward. Co-design methods instead optimize hardware and control variables together. A representative example is the NSGA-II-based co-design of an interleaved converter and its controllers in [4]. For aircraft conversion specifically, NSGA-II has recently been applied to the 270 V-to-28 V LLC design problem [16]. These studies make clear that Pareto structure is intrinsic to high-performance converter design.

A second development is the use of data-driven surrogate models to reduce the number of expensive circuit, finite-element, or time-domain evaluations. Surrogates are now an established design-automation tool in power electronics [5]. Recent work has combined feasibility classification and probabilistic surrogates for converter parameter optimization under efficiency and thermal constraints [6], while physics-informed learning has been investigated for closed-loop DC–DC converter simulation [7]. The attraction is obvious: once a sufficiently accurate approximation is available, very large design spaces can be screened at low cost. The danger is equally important in safety-critical applications. A surrogate can be locally inaccurate precisely near saturation, stability, or thermal boundaries. Optimizer-generated claims should therefore be verified against the governing physics.

Fractional-order control adds another design degree of freedom. The $PI^\lambda D^\mu$ framework introduced into control engineering by Podlubny [1] allows the frequency response and memory of the controller to be shaped continuously through noninteger orders. Its additional flexibility has motivated numerous converter applications, but that flexibility also enlarges the tuning problem. Treating λ as an optimization variable together with L , C , switching frequency, and feedback gains produces a natural hardware–control co-design problem.

This paper develops such a formulation for an aerospace-relevant three-phase interleaved DC–DC converter. The contribution is deliberately narrower than a claim of hardware certification. The model is a reproducible design-screening environment intended to answer a precise question: *how do transient performance, electrical loss, passive sizing, and fault ride-through compete when fractional controller memory and power-stage parameters are optimized together?* The main contributions are as follows.

1. A six-dimensional co-design vector

$$\boldsymbol{\theta} = (f_s, L, C, \lambda, K_p, K_i)$$

is introduced for a three-phase interleaved 270 V-to-28 V converter. Hardware and controller variables are therefore optimized in the same design loop.

2. Four objectives are considered simultaneously: a nonlinear transient-performance index, a semiconductor/conduction loss estimate, a normalized passive sizing index, and a conservative capacitor-only fault ride-through allowance. Thermal and ripple requirements are imposed as feasibility constraints rather than being hidden inside a scalar cost.
3. A two-stage machine-learning accelerator is used. A classifier rejects configurations likely to be dynamically infeasible, and a neural surrogate predicts the expensive transient objective. The surrogate is explicitly treated as a screening model; all retained trade-off designs are returned to the original nonlinear simulation.
4. The numerical experiment is reproducible: 5000 Latin-hypercube samples are used for model construction, a fixed random seed is specified, 60,000 additional candidates are screened, and the final 68 trade-off designs are independently re-evaluated.

5. The balanced design demonstrates that improvement is not free. It reduces loss, transient cost, and fault sag while slightly reducing the sizing index, but phase-current ripple increases. This trade-off is reported rather than suppressed.

The rest of the paper is organized as follows. Section 2 develops the converter and fractional controller model. Section 3 defines the objectives and constraints. Section 4 describes the surrogate-assisted workflow. Section 5 presents the Pareto and time-domain results, and Section 6 discusses interpretation, limitations, and the distinction between design screening and hardware validation.

2 Converter and fractional control model

2.1 Averaged interleaved converter

We consider a three-phase synchronous interleaved buck stage supplying a regulated 28 V DC bus from a nominal 270 V source. Under continuous-conduction operation and ideal interleaving, the phase currents are close enough that the total inductor current

$$I(t) = \sum_{j=1}^N i_j(t)$$

provides a compact averaged representation. If each phase uses inductance L and series parasitic resistance r_L , the aggregate model is

$$\dot{I}(t) = \frac{N(t)}{L} (d(t)V_h(t) - v(t)) - \frac{r_L}{L} I(t), \quad (1)$$

$$C\dot{v}(t) = I(t) - \frac{v(t)}{R(t)}. \quad (2)$$

Here $d(t) \in [0, 1]$ is the common duty command, $V_h(t)$ the high-voltage bus, $v(t)$ the low-voltage bus, and $R(t)$ the equivalent load. The number of active phases $N(t)$ is normally three. To test short-duration fault tolerance, one phase is removed for a prescribed reconfiguration interval and the spare path is then assumed to restore $N = 3$. This abstraction represents the outer electrical consequence of a redundant architecture without reproducing the detailed fault-detection logic of a dedicated fault-tolerance study.

The nominal operating point is

$$V_h = 270 \text{ V}, \quad v^* = 28 \text{ V}, \quad P_o = 800 \text{ W}.$$

A 270 V distribution level is well aligned with modernized and future aircraft DC architectures reported in the literature [13, 14].

2.2 Fractional PI^λ controller

Let

$$e(t) = v^* - v(t).$$

The duty command is

$$d(t) = \text{sat}_{[d_{\min}, d_{\max}]} \left[d_0 + K_p e(t) + K_i {}_0I_t^\lambda e(t) \right], \quad 0 < \lambda \leq 1, \quad (3)$$

where $d_0 = v^*/V_{h,0}$ and ${}_0I_t^\lambda$ is the Riemann–Liouville fractional integral of order λ . The controller becomes an ordinary PI law at $\lambda = 1$.

For numerical evaluation, a causal convolution approximation is used:

$${}_0I_{t_n}^\lambda e \approx h^\lambda \sum_{j=0}^n w_j^{(\lambda)} e_{n-j}, \quad w_0^{(\lambda)} = 1, \quad w_j^{(\lambda)} = w_{j-1}^{(\lambda)} \frac{j-1+\lambda}{j}. \quad (4)$$

The approximation is adequate for the design-screening purpose of the present study and retains the dependence of the controller on the error history.

2.3 Test sequence

Every dynamic design evaluation uses the same 20 ms sequence with a 50 μ s integration step:

- $0 \leq t < 5$ ms: nominal 20 A load;
- $t = 5$ ms: load changes to 28 A;
- $9 \text{ ms} \leq t < 9.3$ ms: one phase is unavailable, so $N = 2$;
- $t = 12$ ms: the high-voltage bus changes from 270 V to 235 V.

This single trajectory forces a candidate to face load, fault-reconfiguration, and line disturbances rather than being tuned to one step response.

3 Multiobjective formulation

3.1 Decision set

The decision vector is restricted to the compact hyper-rectangle

$$\mathcal{D} = [150, 500] \times [20, 80] \times [400, 1400] \times [0.65, 1] \times [3 \times 10^{-4}, 3 \times 10^{-3}] \times [0.05, 0.8], \quad (5)$$

where the first three coordinates are expressed in kHz, μ H, and μ F, respectively. These bounds define a study region rather than manufacturing limits for a specific qualified component family.

3.2 Transient objective

The dynamic objective combines the integral time-weighted error with the extreme voltage excursion:

$$J_{\text{tr}} = \frac{1}{10^{-4}} \int_0^T t |e(t)| dt + 0.3 \Delta v_{\text{over}} + 0.5 \Delta v_{\text{under}}, \quad (6)$$

where

$$\Delta v_{\text{over}} = \max(0, \max_t v(t) - v^*), \quad \Delta v_{\text{under}} = \max(0, v^* - \min_t v(t)).$$

The integral term penalizes sustained deviations while the two extremal terms prevent a low integral error from concealing an unacceptable voltage excursion.

3.3 Loss and thermal model

For design-space screening, total electrical loss is approximated by

$$P_\Sigma = P_{\text{cond}} + P_{\text{sw}} + P_g + P_{\Delta i}, \quad (7)$$

with

$$P_{\text{cond}} = NI_{\text{ph}}^2(R_{\text{DS,on}} + r_L), \quad (8)$$

$$P_{\text{sw}} = N \frac{1}{2} V_h I_{\text{ph}} (t_r + t_f) f_s, \quad (9)$$

$$P_g = 2N Q_g V_g f_s, \quad (10)$$

$$P_{\Delta i} = N \left(\frac{\Delta i_L}{\sqrt{12}} \right)^2 (R_{\text{DS,on}} + r_L). \quad (11)$$

The ripple approximation is

$$\Delta i_L \approx \frac{(V_h - v^*)D}{L f_s}, \quad D = \frac{v^*}{V_h}. \quad (12)$$

The numerical screening uses

$$R_{\text{DS,on}} = 20 \text{ m}\Omega, \quad r_L = 18 \text{ m}\Omega, \quad t_r + t_f = 12 \text{ ns}, \quad Q_g = 35 \text{ nC}, \quad V_g = 15 \text{ V}.$$

These values are representative device-level parameters; they are not claimed as a model of a particular commercial switch.

A first-order junction-temperature screen is

$$T_j = T_a + R_{\theta,\text{eff}} P_{\text{semi,ph}}, \quad (13)$$

with $T_a = 70^\circ\text{C}$ and $R_{\theta,\text{eff}} = 2 \text{ K/W}$. Detailed thermal networks, mission-dependent cycling, and package-specific heat spreading are intentionally outside the present model. Their importance for real high-power-density converters is well established [9, 8].

3.4 Passive sizing and fault ride-through

Instead of claiming an absolute component mass without a selected inductor core and capacitor technology, we define a transparent normalized sizing index

$$S_{\text{pass}} = 0.60 \frac{L}{40 \text{ }\mu\text{H}} + 0.40 \frac{C}{600 \text{ }\mu\text{F}}. \quad (14)$$

Thus $S_{\text{pass}} = 1$ at the baseline design. The index is only a co-design proxy; it should be replaced by component-database mass or volume in hardware-oriented optimization.

A conservative capacitor-only ride-through allowance follows from

$$C \frac{dv}{dt} = -\Delta I$$

during a temporary current deficit $\Delta I > 0$. Allowing a voltage drop ΔV_{adm} gives

$$t_{\text{ride}} = \frac{C \Delta V_{\text{adm}}}{\Delta I}. \quad (15)$$

We use $\Delta V_{\text{adm}} = 2 \text{ V}$ and $\Delta I = 8 \text{ A}$.

Proposition 3.1 (Conservative ride-through bound). *If the converter-side current deficit satisfies $0 \leq \Delta I(t) \leq \Delta I_{\text{max}}$ on $[t_f, t_f + \tau]$ and no other source supports the output capacitor, then*

$$v(t_f + \tau) \geq v(t_f) - \frac{\Delta I_{\text{max}}}{C} \tau.$$

Consequently, the sufficient condition

$$\tau \leq \frac{C \Delta V_{\text{adm}}}{\Delta I_{\text{max}}}$$

guarantees $v(t_f + \tau) \geq v(t_f) - \Delta V_{\text{adm}}$.

Proof. Integrating $C\dot{v} = -\Delta I(t)$ from t_f to $t_f + \tau$ gives

$$v(t_f + \tau) = v(t_f) - \frac{1}{C} \int_{t_f}^{t_f + \tau} \Delta I(s) ds.$$

The assumed bound on ΔI yields the stated inequality. $\square$

3.5 Optimization problem

The four-objective problem is

$$\min_{\boldsymbol{\theta} \in \mathcal{F}} (J_{\text{tr}}(\boldsymbol{\theta}), P_{\Sigma}(\boldsymbol{\theta}), S_{\text{pass}}(\boldsymbol{\theta}), -t_{\text{ride}}(\boldsymbol{\theta})), \quad (16)$$

where the feasible set requires

$$T_j < 125^\circ\text{C}, \quad \Delta i_L < 8\text{ A}, \quad 20\text{ V} < v(t) < 38\text{ V} \quad \text{throughout the test sequence.} \quad (17)$$

Proposition 3.2 (Existence of Pareto-optimal designs). *Assume the time-domain solution of (1)–(3) depends continuously on $\boldsymbol{\theta}$ for every dynamically feasible design and that the feasibility inequalities in (17) are closed after replacing strict bounds by engineering margins. If $\mathcal{F}$ is nonempty, then the multiobjective problem (16) possesses at least one Pareto-optimal point.*

Proof. The design set (5) is compact. Under the assumptions, the margin-defined feasible subset is closed and hence compact. Each objective is continuous on that set. Therefore the image of $\mathcal{F}$ under the vector objective is compact in $\mathbb{R}^4$. A compact nonempty objective image contains minimal elements under the componentwise order, which correspond to Pareto-optimal designs. $\square$

4 Surrogate-assisted design workflow

The surrogate is not permitted to certify a final design. Its role is to reject clearly unpromising regions and to rank a much larger pool than would otherwise be convenient. The workflow is shown in Fig. 1.




Figure 1: Physics-first surrogate-assisted co-design workflow. Machine learning screens the search space; final results are always recomputed with the nonlinear model.

4.1 Training database and feasibility classifier

A maximin-style Latin-hypercube sample of 5000 points is generated in (5) using random seed 20260913. Of these designs, 3595 satisfy the dynamic, thermal, and ripple constraints. A tree-ensemble classifier is trained on 80% of the data and tested on the remaining 20%, stratified by feasibility. The confusion matrix is

$$\begin{pmatrix} 274 & 7 \\ 29 & 690 \end{pmatrix},$$

with infeasible and feasible classes ordered as rows and predictions as columns. Thus

$$\text{accuracy} = 96.4\%, \quad \text{precision} = 99.0\%, \quad \text{recall} = 96.0\%.$$

The high precision is useful here because false feasible predictions waste expensive full-model verification.

4.2 Transient surrogate

The expensive response-derived objective J_{tr} is approximated by a multilayer perceptron with hidden widths 256–128–64, standardized inputs, logarithmically transformed positive target, and early stopping. On the held-out stable-design test set, the model gives

$$R^2 = 0.495, \quad \text{MAPE} = 7.18\%, \quad \text{median APE} = 3.56\%.$$

The moderate R^2 is not concealed: saturation and extremum terms in (6) create sharp local changes that are difficult for a smooth surrogate. The model is therefore not used as an acceptance test. It is only a ranking device.

Figure 2 shows a representative subset of test predictions.




Figure 2: Held-out transient-objective predictions. The dispersion near the extremes motivates full-model verification of all final designs.

4.3 Large-pool screening

A second set of 60,000 uniformly distributed candidates is generated with random seed 20260911. A design is retained for multiobjective ranking only when the classifier assigns feasibility probability at least 0.90 and the explicit thermal and ripple expressions also satisfy (17). This leaves 26,209 candidates.

Rather than collapse the four objectives into one fixed weighted sum, 2000 weight vectors are drawn from a symmetric Dirichlet distribution. For each weight vector, the normalized weighted objective is minimized. The union is reduced by dominance, producing 68 trade-off designs. Every one of these 68 designs is then simulated again with (1)–(3); the reported tables and plots use those verified values.

5 Numerical results

5.1 Baseline and representative Pareto designs

The conventional reference uses

$$(f_s, L, C, \lambda, K_p, K_i) = (300 \text{ kHz}, 40 \text{ }\mu\text{H}, 600 \text{ }\mu\text{F}, 1, 0.001, 0.2).$$

Table 1 compares this baseline with three full-model-verified trade-off designs. “Balanced” is selected because it improves all four optimization objectives relative to the baseline without approaching the imposed ripple limit.

Table 1: Design variables of representative verified solutions.

Design	f_s (kHz)	L (μ H)	C (μ F)	λ	K_p	K_i
Baseline	300.00	40.00	600.00	1.0000	0.001000	0.2000
Low-loss/ride	150.40	38.83	1358.16	0.8283	0.000611	0.5292
Compact	210.11	37.36	427.65	0.9753	0.000691	0.7676
Balanced	246.23	22.53	960.91	0.9042	0.000379	0.4679

Table 2: Full-model performance of representative designs.

Design	J_{tr}	P_Σ (W)	S_{pass}	t_{ride} (ms)	fault sag (V)	Δi_L (A)
Baseline	6.446	24.525	1.000	0.150	0.198	2.091
Low-loss/ride	4.604	17.402	1.488	0.340	0.185	4.297
Compact	5.324	20.220	0.845	0.107	0.199	3.197
Balanced	5.556	22.069	0.979	0.240	0.126	4.523

The table exposes the geometry of the design problem. The low-loss/ride solution has the smallest loss and strongest capacitor support among the listed points, but its passive index is almost 49% above baseline. The compact design reduces passive size substantially but sacrifices ride-through time. The balanced solution improves the baseline in all four formal objectives, yet its ripple is more than twice the baseline value. Multiobjective optimization does not remove the trade-off; it makes the trade-off explicit.

5.2 Pareto structure

Figure 3 shows the verified 68-design set projected onto two two-dimensional planes. Projections should not be confused with the full four-dimensional dominance relation, but they reveal the main compromises.




Figure 3: Verified trade-off set. Squares denote the baseline; triangles denote the balanced solution.

5.3 Time-domain comparison

Figure 4 compares the baseline and balanced designs under the common test sequence. The balanced design does not simply use a more aggressive controller. Its lower switching frequency, smaller phase inductance, larger capacitance, and noninteger integral order act together.




Figure 4: Output-bus response under the load step, temporary phase loss, and input-bus sag.

Relative to the baseline, the balanced design changes the principal metrics as follows:

$$\begin{array}{ll}
 J_{\text{tr}} : 6.446 \rightarrow 5.556 & (-13.8\%), \\
 \text{IAE} : 2.287 \times 10^{-2} \rightarrow 1.613 \times 10^{-2} & (-29.5\%), \\
 P_{\Sigma} : 24.525 \text{ W} \rightarrow 22.069 \text{ W} & (-10.0\%), \\
 t_{\text{ride}} : 0.150 \text{ ms} \rightarrow 0.240 \text{ ms} & (+60.2\%), \\
 S_{\text{pass}} : 1.000 \rightarrow 0.979 & (-2.1\%), \\
 \Delta v_{\text{fault}} : 0.198 \text{ V} \rightarrow 0.126 \text{ V} & (-36.2\%).
 \end{array}$$

The estimated efficiency rises from 97.03% to 97.32%, while the screening junction temperature falls from 83.19 °C to 81.45 °C. The countervailing result is

$$\Delta i_L : 2.09 \text{ A} \rightarrow 4.52 \text{ A},$$

an increase of approximately 116%. It remains below the imposed 8 A constraint, but this is exactly the kind of cost that a single-objective optimization could hide.

5.4 Effect of the fractional order

The Pareto set contains fractional orders throughout the permitted interval rather than concentrating at one fixed value. The balanced design uses $\lambda = 0.9042$, the low-loss/ride design uses $\lambda = 0.8283$, and the compact design returns close to the integer case at $\lambda = 0.9753$. This behavior is consistent with the interpretation of λ as a continuous memory-shaping parameter rather than a universally optimal constant. In the present formulation, the fractional order is useful because it expands the set of attainable hardware–control compromises; the results do not support the stronger claim that a fractional controller dominates an integer controller for every objective.

6 Discussion

6.1 Why the AI layer is deliberately subordinate to physics

The numerical results support a restricted but useful role for machine learning. The feasibility classifier is accurate enough to remove large parts of the design space, while the transient surrogate reduces the number of full dynamic evaluations needed during broad screening. Yet its R^2 value is only moderate. The reason is physically understandable: duty saturation, extremum penalties, and transitions between stable and poorly damped responses introduce nonsmooth behavior in the

target. A surrogate trained only on average regression error can therefore be misleading near the boundaries that matter most to an aerospace designer.

For this reason, the workflow does not report surrogate-generated performance as if it were physics. All 68 retained solutions are returned to the nonlinear averaged model. This verification principle is compatible with the broader movement toward physics-aware and trustworthy learning for converter and power-system dynamics [6, 7].

6.2 Aerospace interpretation

Recent aircraft-power research emphasizes that reliability, safety, mass, efficiency, and fast reconfiguration compete at system level [11, 12, 14]. The present results show the same competition at converter level. Larger capacitance is valuable for ride-through, but an optimizer that only maximizes hold-up time pushes toward bulky designs. Higher switching frequency can reduce ripple for fixed L yet raises switching and gate-drive losses. Lower inductance saves the passive sizing index but increases ripple. Controller memory changes the transient response without requiring a discrete controller-family decision.

This paper should therefore be read as a *co-design map*, not as a certified aerospace converter. Certification-oriented design would require actual component databases, device nonlinearities, electromagnetic compatibility models, insulation and partial-discharge constraints, thermal cycling under flight mission profiles, and fault-coverage evidence. The literature on future aircraft EPS architecture and converter reliability makes clear that these layers cannot be collapsed into one averaged simulation [10, 14].

6.3 Relationship to existing multiobjective converter design

The methodology differs from conventional fixed-controller hardware optimization in two ways. First, the fractional order and controller gains are part of the design vector, so the optimizer can exchange dynamic memory against passive and switching choices. Second, surrogate predictions are not accepted at face value; the final trade-off set is reconstructed using full model evaluations. This is complementary to NSGA-II converter co-design [4, 16] and to recent stochastic-surrogate converter optimization [6].

The use of a normalized passive sizing index is intentionally conservative in its claims. It avoids inventing a mass value without choosing specific cores, winding technologies, capacitors, and cooling hardware. In a product-level study, (14) should be replaced by a database-driven mass, volume, cost, or specific-power model.

6.4 Limitations

Several limitations define the boundary of the conclusions.

First, the plant is an averaged continuous-conduction model. Switching-node waveforms, dead time, parasitic ringing, common-mode currents, EMI filters, and magnetic saturation are not simulated. Second, the thermal model is a screening relation rather than a multidimensional electrothermal model. Third, the semiconductor parameters are representative values, not a named qualified device. Fourth, the fault event is represented by a temporary reduction in the number of active phases; the diagnostic and reconfiguration hardware is not modeled. Fifth, the passive sizing index is dimensionless. Finally, no hardware-in-the-loop or laboratory experiment is claimed.

These limitations are not cosmetic. They imply that the numerical values are evidence for the *design methodology* and the identified trade-offs, not proof that a specific aerospace converter meets a certification standard.

7 Conclusion

A surrogate-assisted multiobjective co-design framework has been developed for an aerospace-relevant interleaved DC–DC converter with fractional-order voltage control. The method treats switching frequency, inductance, capacitance, fractional order, and controller gains as a single design vector and evaluates four competing objectives: transient quality, loss, passive sizing, and fault ride-through.

The machine-learning layer accelerates screening but does not replace the physics model. From 5000 physics-generated training designs, the feasibility classifier achieved 96.4% test accuracy. The transient surrogate was less accurate near dynamic boundaries, with 7.18% mean absolute percentage error; accordingly, all retained designs were re-simulated before being reported. The final verified set shows that no single parameter choice dominates every criterion.

A balanced solution with $\lambda = 0.9042$ reduced the transient objective by 13.8%, estimated converter loss by 10.0%, and fault sag by 36.2%, while increasing conservative capacitor ride-through by 60.2%. The normalized passive sizing index decreased slightly. These gains were purchased with substantially larger phase-current ripple, although the ripple remained feasible. This result captures the central point of the paper: robust aerospace converter design is a trade-off problem, and fractional control is most useful when it is optimized together with the hardware rather than appended after hardware selection.

Future work should replace screening proxies with component-database and electrothermal models, include conducted/radiated EMI metrics, quantify uncertainty in device and mission parameters, and validate selected designs through switching simulation, hardware-in-the-loop testing, and an experimental converter prototype.

Data and computational reproducibility

No external experimental data are used. The reported design study is generated from the equations in this paper. The training set uses 5000 Latin-hypercube points with seed 20260913. The large-pool screening uses 60,000 candidate points with seed 20260911. The time-domain integration step is 50 μ s, and each trajectory spans 20 ms. All final design metrics are recomputed from the original nonlinear averaged model rather than copied from surrogate predictions.

Declarations

Competing interests. The authors declare no competing interests.

Funding. To be completed by the authors before submission.

Author contributions. To be finalized and approved by both authors before submission.

Use of AI/ML in the scientific method. Machine learning is part of the proposed computational methodology for feasibility screening and transient-objective approximation. Final numerical claims are based on re-evaluation with the physics model.

References

- [1] I. Podlubny, Fractional-order systems and $PI^\lambda D^\mu$ controllers, *IEEE Transactions on Automatic Control* 44(1) (1999) 208–214. doi:10.1109/9.739144.
- [2] C. A. Monje, Y. Chen, B. M. Vinagre, D. Xue and V. Feliu, *Fractional-order Systems and Controls: Fundamentals and Applications*, Springer, London, 2010.
- [3] K. Deb, A. Pratap, S. Agarwal and T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computation* 6(2) (2002) 182–197. doi:10.1109/4235.996017.

- [4] D.-D. Tran, S. Chakraborty, Y. Lan, M. El Baghdadi and O. Hegazy, NSGA-II-based codesign optimization for power conversion and controller stages of interleaved boost converters in electric vehicle drivetrains, *Energies* 13(19) (2020) 5167. doi:10.3390/en13195167.
- [5] Y. Zhang, Y. Xu and F. Blaabjerg, Surrogate models for power electronic systems applying machine learning techniques, in: *Control of Power Electronic Converters and Systems*, Vol. 4, Elsevier, 2024, pp. 333–352. doi:10.1016/B978-0-323-85622-5.00002-X.
- [6] A. Mahajan, S. Chaturvedi, S. Das, W. Su and V.-H. Bui, Optimal parameter design for power electronic converters using a probabilistic learning-based stochastic surrogate model, *Next Energy* 9 (2025) 100464. doi:10.1016/j.nxener.2025.100464.
- [7] M.-A. Coulombe, M. Berger and A. Lesage-Landry, Simulation of a closed-loop DC–DC converter using a physics-informed neural network-based model, *Electric Power Systems Research* 252 (2026) 112408. doi:10.1016/j.epsr.2025.112408.
- [8] M. Buffolo, D. Favero, A. Marcuzzi, C. De Santi, G. Meneghesso, E. Zanoni and M. Meneghini, Review and outlook on GaN and SiC power devices: industrial state-of-the-art, applications, and perspectives, *IEEE Transactions on Electron Devices* 71(3) (2024) 1344–1355. doi:10.1109/TED.2023.3346369.
- [9] S. M. I. Rahman et al., Emerging trends and challenges in thermal management of power electronic converters: a state of the art review, *IEEE Access* 12 (2024) 50633–50672. doi:10.1109/ACCESS.2024.3385429.
- [10] A. Verma et al., Reliability analysis of multilevel and matrix converters used in more electric aircraft, *IET Electrical Systems in Transportation* (2023). doi:10.1049/els2.12078.
- [11] G. Martinez, F. Rodríguez, I. Sanchez-Guardamino, M. Urnieta and J. M. Echeverría, Improving the reliability of a future more electric aircraft DC power system through a novel topology and software combination, *International Journal of Electrical Power & Energy Systems* 157 (2024) 109795. doi:10.1016/j.ijepes.2024.109795.
- [12] X. Wang, J. Atkin and S. Bozhko, Fault-tolerant hierarchical energy management system for an electrical power system on more-electric aircraft, *Applied Energy* 379 (2025) 124955. doi:10.1016/j.apenergy.2024.124955.
- [13] I. Hakyemez, O. K. Keles, O. Kalenderli and M. Bagriyanik, Conceptual design of modernized 270 VDC electrical power distribution system for Hercules C-130 aircraft, *IEEE Access* 12 (2024) 112053–112061. doi:10.1109/ACCESS.2024.3443189.
- [14] A. Reindl and F. Van der Linden, Architectural design considerations for electric power systems in future (more) electric aircraft, *Engineering Proceedings* 133(1) (2026) 83. doi:10.3390/engproc2026133083.
- [15] L. Setlak, R. Kowalik, A. Gębura et al., Dynamic stability analysis of the aircraft electrical power system in the more electric aircraft concept, *Scientific Reports* 14 (2024) 25521. doi:10.1038/s41598-024-74762-1.
- [16] T. Y. Kiymaz, M. Taşkıran, A. H. Obdan and K. N. Bekiroğlu, NSGA-II based optimization of LLC resonant converter for 270 V to 28 V DC–DC conversion in more electric aircraft, in: *2025 7th International Congress on Human-Computer Interaction, Optimization and Robotic Applications (ICHORA)*, 2025, pp. 1–5. doi:10.1109/ICHORA65333.2025.11016978.